

Canonical Divergence Analysis

Hoang-Vu Nguyen[°]

Jilles Vreeken[°]

Abstract

We aim to analyze the relation between two random vectors that may potentially have both different number of attributes as well as realizations, and which may even not have a joint distribution. This problem arises in many practical domains, including biology and architecture. Existing techniques assume the vectors to have the same domain or to be jointly distributed, and hence are not applicable. To address this, we propose Canonical Divergence Analysis (CDA). We introduce three instantiations, each of which permits practical implementation. Extensive empirical evaluation shows the potential of our method.

1 Introduction

In many application domains we want to analyse the relation between data sets (views) of different sizes, dimensionality, and often also without sample-to-sample correspondence (i.e. no joint distribution). For instance, in metabolomics the goal is to study (dis)similarities of metabolic profiles collected from different populations, possibly different species, e.g. humans and mice [29,43]. The results enable combining profiles from different organisms for joint analysis, leading to a more comprehensive view of biological properties. For energy efficient architecture, on the other hand, it is useful to know the impact of various energy and climate indicators over different buildings [44]. As the data sets may be on different domains, e.g. different sets of indicators with different cardinality recorded by heterogeneous technologies, to facilitate this it is necessary to find mappings that reveal similarities. In general, this problem also appears in e.g. comparing subspace clusters [41], finding structural similarities between different communities in a social network [42], and matching schemas between (legacy) databases [46].

Formally, the problem can be cast as follow. Consider two random vectors, i.e. two sets of random variables. The first vector $\mathbf{X} \in \mathbb{R}^m$ has n realizations $\mathbf{x}_1, \dots, \mathbf{x}_n$. The second $\mathbf{Y} \in \mathbb{R}^l$ has k realizations $\mathbf{y}_1, \dots, \mathbf{y}_k$. In general, $m \neq l$ and $n \neq k$. That is, \mathbf{X} and \mathbf{Y} may differ on number of attributes *as well as* on number of realizations. Further, there is no sample-to-sample correspondence between them, i.e. they do not have a joint distribution. Our goal is to find mappings of \mathbf{X} and \mathbf{Y} maximizing their similarity, or equivalently minimizing their divergence.

Surprisingly, existing work does not address this general problem. Multi-view learning [40], e.g. canonical correlation analysis [17], requires that \mathbf{X} and \mathbf{Y} are jointly distributed; that is, they must have a sample-to-sample correspondence. Transfer learning techniques [31] on the other hand in general require that \mathbf{X} and \mathbf{Y} are from the same domain, i.e. $m = l$. Sample-to-sample correspondence analysis [43] assumes that each realization of \mathbf{X} is paired with exactly one realization of \mathbf{Y} . Thus, we can use neither to analyze data from e.g. different species or buildings.

To allow for meaningful analysis of data sets of different sizes, dimensionality, and possibly without a joint distribution, we propose Canonical Divergence Analysis (CDA). We instantiate it with three different divergence measures, leading to practical solutions. Further, we introduce two alternate formulations of CDA, each with specific strengths; in particular, one for global optimization and one for fast optimization. Extensive experiments on a wide range of tasks demonstrate the potential of our approach.

2 Formulation

We denote $r = \min\{n, k\}$. We also write u for a $m \times 1$ vector and v for a $l \times 1$ vector. We define CDA as follow.

DEFINITION 1. *Canonical Divergence Analysis (CDA) iteratively solves for $\{(u^i, v^i)\}_{i=1}^r$ where $u^i \in \mathbb{R}^m$ and $v^i \in \mathbb{R}^l$ at the i^{th} iteration are the solutions of*

$$\arg \min_{\substack{u^T u = v^T v = 1 \\ u^T u^j = 0, \forall j \in [1, i-1] \\ v^T v^j = 0, \forall j \in [1, i-1]}} \text{diff}(p(u^T \mathbf{X}), q(v^T \mathbf{Y}))$$

with $p(\cdot)$ and $q(\cdot)$ being pdfs, diff being a divergence measure of pdfs, and $\beta \neq 0$ being the scaling factor to bring $u^T \mathbf{X}$ and $v^T \mathbf{Y}$ to the same domain.

There are some points to note. First, the *orthonormal constraints* here prevent the canonical vectors from *degenerating*. In addition, the orthogonal constraints, i.e. the second and third constraints, are not applied when solving for (u^1, v^1) ; they are used when learning $\{(u^i, v^i)\}_{i=2}^r$.

Second, one could compare $u^T \mathbf{X}$ and $v^T \mathbf{Y}$ directly. They may however not be on the same domain, potentially causing biases. For fair comparison, we use the scaling factor β , which is also learned.

[°]Max Planck Institute for Informatics and Saarland University, Germany. Email: {hnguyen, jilles}@mpi-inf.mpg.de

Third, in CDA we quantify the (dis)similarity of the transformed data $u^T \mathbf{X}$ and $\beta v^T \mathbf{Y}$ through the divergence of their pdfs. This has intuition from statistics. In particular, the Mallows distance [34] of two univariate/multivariate random variables with no joint distribution is equal to the minimum of their expected difference – defined by a given metric – taken over *all* possible joint distributions that preserve their marginal pdfs. Mallows distance is more commonly known as the Earth Mover’s Distance [22], a divergence measure. In other words, Mallows distance tells us that analyzing the relationship between two random variables, taking into account all permissible joint distributions, reduces to computing the divergence of their distributions.

Fourth, CDA has connections to the normalized information distance (NID) [10] from information theory. Loosely speaking, NID is based on the minimal Kolmogorov complexity of expressing \mathbf{X} given \mathbf{Y} , or vice versa. That is, NID considers the length of the shortest algorithm that can generate one data set given the other. Mapping to CDA, this generation is done through finding u , v , and β ; the complexity is expressed by $\text{diff}(p(u^T \mathbf{X}), q(\beta v^T \mathbf{Y}))$.

Furthermore, CDA could be considered as a generalization of transfer component analysis (TCA) [30]. In short, TCA assumes that \mathbf{X} and \mathbf{Y} are from the *same* domain, i.e. $m = l$. It learns a *common* mapping ϕ that minimizes $\text{diff}(p(\phi(\mathbf{X})), q(\phi(\mathbf{Y})))$. CDA in turn allows $m \neq l$ and hence seeks for two different mappings. In other words, CDA considers a more general problem setting.

Finally, CDA focuses on linear mappings, which facilitates interpretation and post analysis [9, 18]. It however does not limit CDA to linear relations; the complexity of the latent similarity discovered is determined by diff . That is, if diff is non-linear, CDA will be able to discover non-linear relationships. Further non-linear mappings can be achieved by searching for kernel transformations [16] or performing deep learning [3]. As this is the first work on CDA, we hence focus on the problem formulation and how to instantiate it for practical analysis – we postpone extending CDA to non-linear mappings to future work, but do note that our experiments show that with our current formulation far more complex relations than linear can be discovered.

Having introduced CDA, next we review related work.

3 Related Work

Multi-view learning [40] searches for transformations of \mathbf{X} and \mathbf{Y} of different domains. For instance, CCA [17] and its extensions [3, 9, 16, 45] study the relationships between \mathbf{X} and \mathbf{Y} by searching for mappings maximizing their correlation. For correlation analysis, the joint distribution of \mathbf{X} and \mathbf{Y} must exist. This means that $n = k$ and every realization of \mathbf{X} corresponds to a unique realization of \mathbf{Y} (and vice versa). Therewith CCA and multi-view learning in general address a different, more specific problem.

CDA is related to multi-target regression, e.g. [1], in the sense that one can perceive \mathbf{X} as the independent variables and \mathbf{Y} as the targets. This setting however also requires \mathbf{X} and \mathbf{Y} to have a sample-to-sample correspondence.

Sample-to-sample correspondence analysis [43] assumes that each realization of \mathbf{X} is paired with exactly one realization of \mathbf{Y} , but not necessarily the reverse. It first searches for mappings of \mathbf{X} and \mathbf{Y} to bring them to a common domain. Then, it identifies a permutation of the realizations of \mathbf{Y} after mapping that minimizes the difference to the mapped realizations of \mathbf{X} . By assuming sample correspondence, it is not suited to analysing data taken from e.g. different species or different buildings. CDA in contrast does not make this assumption and hence is more flexible.

In transfer learning [31], \mathbf{X} and \mathbf{Y} can stand for the variables in the source and target domains. The goal is to learn transformations of \mathbf{X} and \mathbf{Y} that are then used to boost performance of a learning task, e.g. classification or clustering, in the target domain. For instance, transductive transfer learning [4] identifies mappings minimizing reconstruction error to the respective original space and classification error, i.e. its mappings is tailored to classification. Transfer component analysis [30], assuming that \mathbf{X} and \mathbf{Y} are from the same domain, learns a common mapping where distributions of the mapped \mathbf{X} and \mathbf{Y} are similar. Likewise, self-taught learning [35], structural correspondence learning [7], and domain adaptation with coupled subspaces [6] assume that \mathbf{X} and \mathbf{Y} are from the same domain. Self-taught clustering [11] searches for mappings of \mathbf{X} and \mathbf{Y} that bring them to the same domain with the aim to help clustering in the target domain. It however works on discrete data and assumes that a common feature space of \mathbf{X} and \mathbf{Y} is given. A comprehensive review of transfer learning literature can be found in [31]. CDA on the other hand is task-independent and works directly with real-valued data. Given \mathbf{X} and \mathbf{Y} from *different* domains and *without* prior knowledge about their common space, we learn mappings that transform them to the same domain to facilitate joint analysis.

Next, we solve CDA by introducing different instantiations of divergence measure diff . In Section 5 we discuss how to learn β during the optimization process. In addition, in Section 6 we propose two alternative ways to formulate the CDA problem which enrich our options towards tackling it. Finally, we discuss our choice of the optimizer in Section 7.

4 Divergence Measure diff

With diff we quantify the divergence between two univariate pdfs $p(u^T \mathbf{X})$ and $q(\beta v^T \mathbf{Y})$. If we want to optimize diff it helps if it is differentiable. We therefore focus on three such divergence measures. The first is an extension of Mallows distance [34]. The second and third are based on kernel density estimation (KDE) [39]. Note that we are

not constrained by these three measures; any differentiable divergence measure, e.g. Maximum Mean Discrepancy [8], can be straightforwardly plugged into our solution.

Before going into the details, we introduce some notation. In general, it only makes sense to compare $p(u^T \mathbf{X})$ and $q(\beta v^T \mathbf{Y})$ if they are defined on the same univariate random variable. We denote such a variable as Z . Hence, $p(u^T \mathbf{X})$ is equivalent to $p(Z)$ and $q(\beta v^T \mathbf{Y})$ is equivalent to $q(Z)$. Under either notation we write the i.i.d. realizations of $p(\cdot)$ as x_1, \dots, x_n where $x_i = u^T \mathbf{x}_i$. Analogously, we write the i.i.d. realizations of $q(\cdot)$ as y_1, \dots, y_k where $y_j = \beta v^T \mathbf{y}_j$.

4.1 Extended Mallows Distance For our first measure we extend the well-known Mallows distance [34]. For exposition, we denote $u^T \mathbf{X}$ as X and $\beta v^T \mathbf{Y}$ as Y . Our goal is to measure $\text{diff}(p(X), q(Y))$. Let \mathcal{F} be the set of joint pdfs $f(X, Y)$ of X and Y such that the marginal distributions of f on X and Y are $p(X)$ and $q(Y)$, respectively. That is, $\int_{\text{dom}(Y)} f(x, y) dy = p(x) \quad \forall x \in \text{dom}(X)$, and $\int_{\text{dom}(X)} f(x, y) dx = q(y) \quad \forall y \in \text{dom}(Y)$.

Given $t \geq 1$, the t -th Mallows distance between $p(X)$ and $q(Y)$ is [22]

$$MA_t(p(X), q(Y)) = \left(\min_{f \in \mathcal{F}} \int |x - y|^t f(x, y) dx dy \right)^{1/t}$$

If $p(X)$ and $q(Y)$ have the same number of realizations ($n = k$), we have $\text{Mallows}_t(p(X), q(Y)) = \left(\frac{1}{n} \min_{\{j_1, \dots, j_n\}} \sum_{i=1}^n |x_i - y_{j_i}|^t \right)^{1/t}$ where $\{j_1, \dots, j_n\}$ is a permutation of $\{1, \dots, n\}$. If $n \neq k$, a possible solution is to replicate each realization so that both distributions have $n \times k$ realizations each. This practice yields the same Mallows distance as it preserves the empirical distributions.

Calculating Mallows distance on empirical data requires us to select the best permutation, which is expensive [22]. To avoid this, we propose a modified version of Mallows distance. In particular, we set

$$\text{diff}(p(X), q(Y)) = \left(\sum_{f \in \mathcal{F}} \int |x - y|^t f(x, y) dx dy \right)^{1/t}$$

The difference between this version of diff and vanilla Mallows distance is that the latter takes the minimum while we consider the *sum* over all $f \in \mathcal{F}$. This difference is reminiscent to the relationship between Cramér-von Mises and Kolmogorov-Smirnov tests. In particular, the former computes the integral of the (squared) difference between two cumulative distribution functions (cdfs) over the domain of the underlying random variable. The latter in turn takes the maximum (supremum) difference between two cdfs. We prove the following result.

LEMMA 4.1. *Empirically, minimizing $\text{diff}(p(X), q(Y))$ is equivalent to minimizing $\left(\sum_{i=1}^n \sum_{j=1}^k |x_i - y_j|^t \right)^{1/t}$.*

Proof. We postpone the proof to Appendix A.

For simplicity, we use the function $\sum_{i=1}^n \sum_{j=1}^k |x_i - y_j|^t$ for optimization. In this paper, we set $t = 2$ and have

$$\text{diff}(p(X), q(Y)) = \sum_{i=1}^n \sum_{j=1}^k (x_i - y_j)^2.$$

That is, we resemble $\text{diff}(p(X), q(Y))$ as a least-squares cost function. We will refer to CDA using the extended Mallows distance as CDA_M .

Interestingly, CDA_M is closely related to linear CCA following the below lemma by Golub and Zha [14].

LEMMA 4.2. *Assume that \mathbf{X} and \mathbf{Y} are whitened. Let (u^i, v^j) be the i^{th} canonical pair of linear CCA. Then*

$$(u^i, v^i) = \arg \min_{\substack{u^T u = v^T v = 1 \\ u^T u^j = 0, j \in [1, i-1] \\ v^T v^j = 0, j \in [1, i-1]}} \|u^T \mathbf{X} - v^T \mathbf{Y}\|^2.$$

Note that in the context of CCA, $n = k$ and each row of \mathbf{X} is connected to exactly one row of \mathbf{Y} . If we let $x_i = u^T \mathbf{x}_i$ and $y_i = v^T \mathbf{y}_i$ then linear CCA is equivalent to finding (u, v) minimizing $\sum_{i=1}^n (x_i - y_i)^2$. Therefore, CDA_M could be seen as a generalization of linear CCA when n and k may differ from each other, and there is no sample-to-sample correspondence between \mathbf{X} and \mathbf{Y} .

4.2 Quadratic Measure As the pdfs under consideration are univariate, KDE is applicable and reliable for density estimation. The KDE of $p(Z)$ is defined as $\hat{p}(z) = \frac{1}{n \cdot \sigma_{\mathbf{X}}} \sum_{i=1}^n \kappa\left(\frac{z - x_i}{\sigma_{\mathbf{X}}}\right) = \frac{1}{n} \sum_{i=1}^n \kappa_{\sigma_{\mathbf{X}}}(z - x_i)$ where $\kappa(\cdot)$ is a symmetric but not necessarily positive kernel function that integrates to one, $\sigma_{\mathbf{X}} > 0$ is a smoothing parameter called the bandwidth, and $\kappa_{\sigma_{\mathbf{X}}}(z) = \frac{1}{\sigma_{\mathbf{X}}} \kappa\left(\frac{z}{\sigma_{\mathbf{X}}}\right)$. Likewise, the KDE of $q(Z)$ is $\hat{q}(z) = \frac{1}{k \cdot \sigma_{\mathbf{Y}}} \sum_{j=1}^k \tau\left(\frac{z - y_j}{\sigma_{\mathbf{Y}}}\right) = \frac{1}{k} \sum_{j=1}^k \tau_{\sigma_{\mathbf{Y}}}(z - y_j)$ with $\tau(\cdot)$, $\sigma_{\mathbf{Y}}$, and $\tau_{\sigma_{\mathbf{Y}}}(\cdot)$ similarly defined.

We next instantiate $\text{diff}(p(Z), q(Z))$ using one of the quadratic measures [36]. In particular, we set

$$\text{diff}(p(Z), q(Z)) = \int (p(z) - q(z))^2 dz$$

where $p(Z)$ and $q(Z)$ are approximated by KDE. Using

KDE, empirically $\text{diff}(p(Z), q(Z)) =$

$$\begin{aligned} & \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \kappa_{\sigma_{\mathbf{X}}}(x_i - x_j) + \frac{1}{k^2} \sum_{i=1}^k \sum_{j=1}^k \tau_{\sigma_{\mathbf{Y}}}(y_i - y_j) \\ & - \frac{1}{n \cdot k} \sum_{i=1}^n \sum_{j=1}^k \tau_{\sigma_{\mathbf{Y}}}(x_i - y_j) - \frac{1}{n \cdot k} \sum_{i=1}^n \sum_{j=1}^k \kappa_{\sigma_{\mathbf{X}}}(x_i - y_j) \end{aligned}$$

In practice, any kernel functions can be used to instantiate $\kappa(\cdot)$ and $\tau(\cdot)$. In this work, we use the popular Gaussian kernel, i.e. we set $\kappa(z) = \tau(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}}$. Following [9, 15], we set $\sigma_{\mathbf{X}} = \text{median}\{||\mathbf{x}_i - \mathbf{x}_j|| : 1 \leq i, j \leq n\}$ and $\sigma_{\mathbf{Y}} = \text{median}\{||\mathbf{y}_i - \mathbf{y}_j|| : 1 \leq i, j \leq k\}$. Note that more advanced empirical estimation of the above quadratic measure is available, e.g. [20]. We stick to the current estimation due to its simplicity and its wide adoption in practice.

We will refer to CDA using this quadratic measure as CDA_Q .

4.3 Pearson Divergence As the third measure, we instantiate $\text{diff}(p(Z), q(Z))$ by relative Pearson divergence [33], or PE divergence for short. That is, we set $\text{diff}(p(Z), q(Z)) =$

$$\begin{aligned} & \int \left(\frac{1}{2}p(z) + \frac{1}{2}q(z) \right) \left(\frac{p(z)}{\frac{1}{2}p(z) + \frac{1}{2}q(z)} - 1 \right)^2 dz \\ & + \int \left(\frac{1}{2}p(z) + \frac{1}{2}q(z) \right) \left(\frac{q(z)}{\frac{1}{2}p(z) + \frac{1}{2}q(z)} - 1 \right)^2 dz. \end{aligned}$$

We denote the first and second terms of $\text{diff}(p(Z), q(Z))$ as $\text{PE}(p(Z) \parallel q(Z))$ and $\text{PE}(q(Z) \parallel p(Z))$, respectively. Next, we present the computation of $\text{PE}(p(Z) \parallel q(Z))$. The computation of $\text{PE}(q(Z) \parallel p(Z))$ follows similarly.

One could compute $\text{PE}(p(Z) \parallel q(Z))$ by estimating $p(Z)$ and $q(Z)$ using KDE. However, as PE divergence is involved in ratio of pdfs, this would magnify the estimation error. To avoid this issue, we follow [23] and model the ratio $\frac{p(z)}{\frac{1}{2}p(z) + \frac{1}{2}q(z)}$ by the kernel model $g(z; \theta) = \sum_{i=1}^d \theta_i \omega(z, x'_i)$ where $\theta = (\theta_1, \dots, \theta_d)$ are parameters to be learned from empirical data, $\{x'_1, \dots, x'_d\} \subseteq \{x_1, \dots, x_n\}$, and $\omega(z, x'_i)$ is the kernel basis function. The kernel centers $\{x'_1, \dots, x'_d\}$ are drawn randomly from $\{x_1, \dots, x_n\}$ where $d = \min(200, n)$ to increase efficiency [18]. In addition, we employ the widely used Gaussian RBF kernel $\omega(z, x'_i) = e^{-\frac{(z-x'_i)^2}{2\sigma_{\mathbf{X}}^2}}$. The bandwidth $\sigma_{\mathbf{X}}$ is set similarly to the case of the quadratic measure.

Computing θ requires us to minimize a squared loss objective function. Fortunately, this sub-problem has an analytical solution [23]. In particular, we have $\hat{\theta} = (\mathbf{E} + \lambda \mathbf{I}_d)^{-1} \mathbf{e}$ where \mathbf{E} is a $d \times d$ matrix with $\mathbf{E}_{i,j} = \frac{1}{2n} \sum_{r=1}^n \omega(x_r, x'_i) \omega(x_r, x'_j) +$

$\frac{1}{2k} \sum_{r=1}^k \omega(y_r, x'_i) \omega(y_r, x'_j)$, \mathbf{e} is a d dimensional vector with $\mathbf{e}_i = \frac{1}{n} \sum_{j=1}^n \omega(x_j, x'_i)$, \mathbf{I}_d is the $d \times d$ identity matrix, and $\lambda \geq 0$ is a regularization parameter chosen by cross-validation.

Therefore, we obtain $\hat{g}(z) = \sum_{i=1}^d \hat{\theta}_i \omega(z, x'_i)$. Empirically, $\text{PE}(p(Z) \parallel q(Z)) =$

$$-\frac{1}{4n} \sum_{i=1}^n \hat{g}(x_i)^2 - \frac{1}{4k} \sum_{j=1}^k \hat{g}(y_j)^2 + \frac{1}{n} \sum_{i=1}^n \hat{g}(x_i) - \frac{1}{2}.$$

We will refer to CDA with PE divergence as CDA_P .

5 Setting scaling factor β

To meaningfully compute the divergence between $u^T \mathbf{X}$ and $v^T \mathbf{Y}$ we need to bring them to the same domain. We do so with scaling factor β . Obviously, we can optimize β as part of the CDA optimization problem. Alternatively, reducing the number of variables to optimize, we propose to directly set β at each step of the optimization process, e.g. at each gradient descent of u and v .

As standard in data analysis, before performing CDA we rescale all attributes $\mathbf{X}_1, \dots, \mathbf{X}_m$ of \mathbf{X} and $\mathbf{Y}_1, \dots, \mathbf{Y}_l$ of \mathbf{Y} to the same domain. W.l.o.g., we assume this domain to be $[0, 1]$. Let \bar{m} be the number of non-zero elements of u and \bar{l} be the number of non-zero elements of v .

LEMMA 5.1. *If $\beta = \sqrt{\frac{\bar{m}}{\bar{l}}}$ then $u^T \mathbf{X}$ and $\beta v^T \mathbf{Y}$ have the same domain, which is $[-\sqrt{\bar{m}}, \sqrt{\bar{m}}]$.*

Proof. We postpone the proof to Appendix A.

This implies that we can gain efficiency by setting β to $\sqrt{\frac{\bar{m}}{\bar{l}}}$. Though straightforward, we prove that it is robust to noisy dimensions, as follows.

LEMMA 5.2. *Let (u, v) be a pair of canonical vectors that CDA discovers on \mathbf{X} and \mathbf{Y} , i.e. $\text{diff}(p(u^T \mathbf{X}), q(\beta v^T \mathbf{Y}))$ is small. Given an integer $c > 0$, let $\mathbf{X}' \in \mathbb{R}^{m+c}$ be a random vector created by adding (noisy) attributes $\mathbf{X}_{m+1}, \dots, \mathbf{X}_{m+c}$ to \mathbf{X} . Given an integer $d > 0$, let $\mathbf{Y}' \in \mathbb{R}^{l+d}$ be a random vector created by adding (noisy) attributes $\mathbf{Y}_{l+1}, \dots, \mathbf{Y}_{l+d}$ to \mathbf{Y} . It holds that (u, v) are identifiable when solving CDA on \mathbf{X}' and \mathbf{Y}' .*

Proof. We postpone the proof to Appendix A.

We empirically evaluate this property by testing with data sets that include noisy attributes (see Section 8.1 and Appendix D). We also compare our setting of β to optimizing β ; the results (see Appendix E) show that there is no clear winner while our method is more efficient.

6 Alternate CDA Formulations

With our CDA formulation, we achieve three instantiations that lead to practical solutions. This formulation however is not unique. In this section, we introduce two alternatives. In particular, instead of iteratively learning pairs of canonical vectors, we look at finding multiple pairs simultaneously, which may be good for avoiding local optima. As the second alternative, we replace the hard unit norm constraints $u^T u = v^T v = 1$ by soft reconstruction costs, and hence enable the application of fast unconstrained optimizers. The details are as follow.

6.1 Multi-dimensional CDA In principle, the first alternative is similar to the original formulation, except for that it searches for all pairs of vectors simultaneously instead of iteratively. This may be good for avoiding local optima. Formally, we have:

DEFINITION 2. *Multi-dimensional CDA (MCDA) solves for $\mathbf{U} = (u^1, \dots, u^r) \in \mathbb{R}^{m \times r}$ and $\mathbf{V} = (v^1, \dots, v^r) \in \mathbb{R}^{l \times r}$ where (\mathbf{U}, \mathbf{V}) is the solution of*

$$\arg \min_{\mathbf{U}^T \mathbf{U} = \mathbf{V}^T \mathbf{V} = \mathbf{I}_r} \text{diff} (p(\mathbf{U}^T \mathbf{X}), q(\mathbf{V}^T \mathbf{Y}))$$

with $p(\cdot)$ and $q(\cdot)$ being pdfs, diff being a divergence measure of pdfs, and $\mathbf{\Gamma} = \text{diag}(\beta_1, \dots, \beta_r)$ being the scaling matrix to bring $(u^i)^T \mathbf{X}$ and $(v^i)^T \mathbf{Y}$ to the same domain.

To solve MCDA, we need diff which is suited to multivariate settings and whose analytical form is easy to compute and differentiable. Note that multiplying univariate diff terms does not work as this requires dimensions in projected spaces to be statistically independent. Of the measures we introduced in Section 4, PE divergence meets this requirement and hence we use it for solving MCDA.

6.2 Reconstruction CDA The unit norm constraints of the original CDA problem may lead to slow optimization. To boost efficiency, we relax them using reconstruction costs [5, 21, 28] and hence enable the application of fast unconstrained optimizers. We hence define the second alternative formulation as follow.

DEFINITION 3. *Reconstruction CDA (RCDA) iteratively solves for $\{(u^i, v^i)\}_{i=1}^r$ where $u^i \in \mathbb{R}^m$ and $v^i \in \mathbb{R}^l$ at the i^{th} iteration are the solutions of*

$$\arg \min_{\substack{u^T u^j = 0, \forall j \in [1, i-1] \\ v^T v^j = 0, \forall j \in [1, i-1]}} \left(\frac{\lambda}{n} \sum_{i=1}^n \|u u^T \tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_i\|^2 + \frac{\delta}{k} \sum_{j=1}^k \|v v^T \tilde{\mathbf{y}}_j - \tilde{\mathbf{y}}_j\|^2 + \text{diff} (p(u^T \mathbf{X}), q(\beta v^T \mathbf{Y})) \right)$$

with $\tilde{\mathbf{x}}_i = \mathbf{x}_i - \mu(\mathbf{X}_i)$ being the centered version of \mathbf{x}_i , $\tilde{\mathbf{y}}_j = \mathbf{y}_j - \mu(\mathbf{Y}_j)$ being the centered version of \mathbf{y}_j , $p(\cdot)$ and $q(\cdot)$ being pdfs, diff being a divergence measure of pdfs, $\lambda > 0$ and $\delta > 0$ being weights of reconstruction costs, and $\beta \neq 0$ being the scaling factor to bring $u^T \mathbf{X}$ and $v^T \mathbf{Y}$ to the same domain.

We define Multi-dimensional RCDA (MRCDA) analogously and give the definition in Appendix C. We show that RCDA and the original CDA are related by proving two following properties of the reconstruction costs.

PROPOSITION 6.1. *If \mathbf{X} is whitened, $\frac{\lambda}{n} \sum_{i=1}^n \|u u^T \tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_i\|^2$ is equivalent to $\lambda \|u u^T - \mathbf{I}_m\|_F^2$ where \mathbf{I}_m is the $m \times m$ identity matrix and $\|\cdot\|_F$ denotes the Frobenius norm.*

Proof. We postpone the proof to Appendix A.

PROPOSITION 6.2. *$\lambda \|u u^T - \mathbf{I}_m\|_F^2$ is equivalent to $\lambda (u^T u - 1)^2$ up to an additive constant.*

Proof. We postpone the proof to Appendix A.

Based on Propositions 6.1 and 6.2, we prove the following.

LEMMA 6.1. *If \mathbf{X} and \mathbf{Y} are whitened, $\lambda \uparrow +\infty$, and $\delta \uparrow +\infty$, then 1) RCDA is equivalent to CDA and 2) MRCDA is equivalent to MCDA.*

Proof. We postpone the proof to Appendix A.

With the the unit norm constraints relaxed, we can solve RCDA and MRCDA with fast unconstrained solvers.

7 Optimization

For the original CDA formulation, we try with different optimizers; in particular, projected gradient descent [12], sequential quadratic programming [27], and natural gradient descent [2]. Empirically, we find the last option to yield the best results. Therefore, we use it in our experiment. The update rule of this optimizer is that $u \leftarrow u \times \exp(-t(u^T \nabla_u f - \nabla_u f^T u))$ and $v \leftarrow v \times \exp(-t(v^T \nabla_v f - \nabla_v f^T v))$ where \exp for a matrix denotes the *matrix exponential* and $t > 0$ is the step size. We choose the optimal step size by Nelder-Mead method.

For MCDA, we also use natural gradient descent.

RCDA and MRCDA are unconstrained optimization problems. As optimizer, similar to [21] we use L-BFGS and Conjugate Gradient and present the results with the former. In practice, we see that CDA methods using reconstruction costs are about 5 times faster than their respective counterpart. We provide more details in Appendix F.

8 Experiments

In the following, we represent the data of \mathbf{X} as a $n \times m$ matrix $\mathbf{D}_\mathbf{X}$ where the i -th row of $\mathbf{D}_\mathbf{X}$ corresponds to \mathbf{x}_i . We define the $k \times l$ matrix $\mathbf{D}_\mathbf{Y}$ similarly. We empirically evaluate our methods on four tasks. First, using synthetic data sets we show that they are able to retrieve known relations. Second, using two data sets from UCI Repository [13] we show that they can be used in cross-domain regression and classification. Third, we use them to discover interesting relations between two real-world data sets taken from studies on architecture [38, 44]. Last, we plug our methods into RESCU [24] to discover non-redundant subspace clusters.

We compare to the original CCA [17] and LCCA [18]—a non-linear CCA method; SCL [7] and DAC [7] for transfer learning; SSC [43] for sample-to-sample correspondence detection. We implemented our methods in Matlab and provide the code in the supplementary material.

8.1 Retrieving Canonical Vectors First we evaluate whether our methods can recover associations between data sets where we know the ground truth. To do so, we generate pairs $(\mathbf{D}_\mathbf{X}, \mathbf{D}_\mathbf{Y})$ with relationships between \mathbf{X} and \mathbf{Y} embedded a priori. In particular, we first generate $n = 1000$ samples for $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_m)$ where $m \geq 5$ and $\mathbf{X}_i \sim \mathcal{N}(0, 1)$. Then we generate k samples for $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_l)$ with $l \geq 5$ according to three types of relations.

Linear relations:

$$\begin{aligned} \mathbf{Y}_1 &= \mathbf{X}_1 + 2\mathbf{X}_2 + \epsilon_1 & \mathbf{Y}_2 &= \mathbf{X}_3 + 2\mathbf{X}_4 + \epsilon_2 \\ \mathbf{Y}_3 &= \mathbf{X}_5 + \epsilon_3 & \mathbf{Y}_4 &= \mathbf{X}_2 + \mathbf{X}_5 + \epsilon_4 \\ \mathbf{Y}_{5,\dots,l} &\sim \mathcal{N}(0, 1) & \epsilon_{1,\dots,4} &\sim \mathcal{N}(0, 0.5) \end{aligned}$$

Mixed relations:

$$\begin{aligned} \mathbf{Y}_1 &= \mathbf{X}_1^2 + 2\mathbf{X}_2 + \epsilon_1 & \mathbf{Y}_2 &= \mathbf{X}_3^3 + 2\mathbf{X}_4 + \epsilon_2 \\ \mathbf{Y}_3 &= \mathbf{X}_5 + \epsilon_3 & \mathbf{Y}_4 &= \mathbf{X}_2 + \mathbf{X}_5 + \epsilon_4 \\ \mathbf{Y}_{5,\dots,l} &\sim \mathcal{N}(0, 1) & \epsilon_{1,\dots,4} &\sim \mathcal{N}(0, 0.5) \end{aligned}$$

Non-linear relations:

$$\begin{aligned} \mathbf{Y}_1 &= \mathbf{X}_1^2 + 2\mathbf{X}_2 + \epsilon_1 & \mathbf{Y}_2 &= \mathbf{X}_3^3 + 2\mathbf{X}_4 + \epsilon_2 \\ \mathbf{Y}_3 &= e^{\mathbf{X}_5} + \epsilon_3 & \mathbf{Y}_4 &= \cos(\mathbf{X}_2 + \mathbf{X}_5) + \epsilon_4 \\ \mathbf{Y}_{5,\dots,l} &\sim \mathcal{N}(0, 1) & \epsilon_{1,\dots,3} &\sim \mathcal{N}(0, 0.5), \epsilon_4 \sim \mathcal{N}(0, 0.1) \end{aligned}$$

We refer to the canonical vectors discovered by a method as $\mathbf{U} = (u_1, \dots, u_r)$ and $\mathbf{V} = (v_1, \dots, v_r)$ where $r = \min\{m, l\}$. That is, each column of \mathbf{U} and \mathbf{V} is a canonical vector. Let \mathbf{U}_{gr} and \mathbf{V}_{gr} be the respective ground truth. Similarly to [18], we quantify performance as

$$\frac{1}{2\sqrt{2r}} (\|\mathbf{U}\mathbf{U}^T - \mathbf{U}_{gr}\mathbf{U}_{gr}^T\|_F + \|\mathbf{V}\mathbf{V}^T - \mathbf{V}_{gr}\mathbf{V}_{gr}^T\|_F)$$

where $\|\cdot\|_F$ stands for the Frobenius norm. Intuitively, the smaller the value the better the result.

We experiment with three settings. First, we generate $\mathbf{D}_\mathbf{X}$ and $\mathbf{D}_\mathbf{Y}$ where $n = k$ and their row order is intact. Second, we keep $n = k$ but randomize the row order of $\mathbf{D}_\mathbf{X}$ and $\mathbf{D}_\mathbf{Y}$. Third, we generate data sets where $n = k$; then we randomly remove $\rho \times k$ records of $\mathbf{D}_\mathbf{Y}$ where $\rho \in \{0.05, 0.1, 0.15, 0.2\}$. In all settings, $m = 7$ and $l = 5$, i.e. \mathbf{X} has two noisy dimensions and \mathbf{Y} has one noisy dimension. As SCL and DAC requires that \mathbf{X} and \mathbf{Y} come from the same domain, to test them we skip the 2 noisy dimensions of \mathbf{X} . Regarding our methods, for succinctness we present the outcome on CDA_Q (normal CDA with the quadratic measure), RCDA_M (Reconstruction CDA with the extended Mallows distance), and MRCDA_P (Multi-dimensional Reconstruction CDA with PE divergence). For RCDA_M and MRCDA_P , similar to [21] we set $\lambda = \delta = 0.5$.

We give the results in Table 1. For the first two settings, each value is the average of 10 runs; standard deviation is small and hence skipped. For the last setting, each value is the average over all runs corresponding to all values of ρ tested and 10 runs per value. Going over the results, we make the following observations.

In the first setting, CCA is best on linear relations. RCDA_M and SSC come in second. On mixed and non-linear relations, *all* of our methods give the best results. This is remarkable as they do not use the joint distribution of \mathbf{X} and \mathbf{Y} .

In the second setting where row order is randomized, as expected, the performance of CDA_Q , RCDA_M , and MRCDA_P is not impacted at all. So are transfer learning and sample-to-sample correspondence methods. CCA methods in turn yield not good results, stemming from the fact that they require row order to work properly.

In the third setting, CCA methods are not applicable. Our methods in contrast are relatively stable under the down-sampling of $\mathbf{D}_\mathbf{Y}$; the increase in error stays below 12%, even for $\rho = 20\%$. This corroborates that CDA is suited for analyzing data with neither known association nor the same number of records.

The results so far show RCDA_M to have good performance across different settings. Thus, for the remaining experiments we pick it for exposition.

8.2 Cross-domain Regression and Classification We aim at studying if the mappings identified by RCDA_M are useful for cross-domain learning tasks, e.g. regression and classification. Cross-domain regression and classification are popular forms to assess transfer learning techniques, e.g. SCL [7], DAC [7], and TCA [30]. As CDA could be perceived as a generalization of TCA, it is interesting to study if the mappings identified by our methods, in particular RCDA_M , are also useful for these tasks.

Setting	Type	CDA			CCA	LCCA	SCL	DAC	SSC
		CDA _Q	RCDA _M	MRCDA _P					
$n = k$ <i>row order intact</i>	Linear	0.31	0.30	0.31	0.26	0.45	0.38	0.38	0.30
	Mixed	0.34	0.34	0.35	0.41	0.45	0.48	0.51	0.46
	Non-linear	0.36	0.35	0.35	0.50	0.48	0.52	0.57	0.54
$n = k$ <i>row order randomized</i>	Linear	0.31	0.30	0.31	0.48	0.53	0.38	0.38	0.31
	Mixed	0.34	0.34	0.35	0.62	0.55	0.48	0.51	0.47
	Non-linear	0.36	0.35	0.35	0.69	0.57	0.52	0.57	0.54
$n \neq k$	Linear	0.33	0.31	0.30	<i>n/a</i>	<i>n/a</i>	0.44	0.41	0.35
	Mixed	0.36	0.36	0.38	<i>n/a</i>	<i>n/a</i>	0.49	0.53	0.52
	Non-linear	0.39	0.38	0.39	<i>n/a</i>	<i>n/a</i>	0.53	0.60	0.56

Table 1: [Lower is better] Frobenius norm errors for different settings. All values are obtained after 10 trials. Best values and comparable ones according to the Wilcoxon signed rank test at $\alpha = 5\%$ are in **bold**.

For cross-domain regression, we use the Bike data from UCI Repository. It consists of two data sets with the same attribute set, one recording bike rental statistics per hour (\mathbf{D}_X), one recording per day (\mathbf{D}_Y). The target variable is the number of bikers. Assuming the target values in \mathbf{D}_Y are not available, i.e. \mathbf{D}_Y is unlabeled, the task is to predict the number of bikers of \mathbf{D}_Y using \mathbf{D}_X . Similar to the setup of Pan et al. [30], we first use RCDA_M to learn the mappings of \mathbf{X} and \mathbf{Y} that bring them to the same latent space. We then apply Regularized Least Square Regression (RLSR), training in the latent space using labeled \mathbf{D}_X and some randomly selected records of \mathbf{D}_Y . Finally, we evaluate the performance using the remaining records of \mathbf{D}_Y . We use Average Error Distance as performance metric. Each result shown is the average of 10 runs.

For cross-domain classification, we use the Energy data described in [25]. \mathbf{D}_X and \mathbf{D}_Y contain data collected from the same building in two non-overlapping periods, i.e. we consider the same domain. The class label is whether the given day is a week day or weekend. The testing setup is similar to above except that we use a SVM with linear kernel as classifier, and use Classification Error as metric.

Note that, to enable comparison with SCL [7], DAC [7], and TCA [30], in both experiments we chose \mathbf{D}_X and \mathbf{D}_Y to have the same number of attributes – this is not a requirement of CDA, however. The results are in Figures 1(a) and 1(b). We can see that RCDA_M performs rather well, on par with TCA, suggesting it is a viable technique for cross-domain learning.

8.3 Discovering Novel Relations We use RCDA_M to discover relations between two data sets taken from two nearby office buildings with similar architectural characteristics. These data sets are from architectural studies [38, 44]. The first one (\mathbf{D}_X) has 4421 records and 15 attributes; some of which are *electricity consumption*, *amount of water con-*

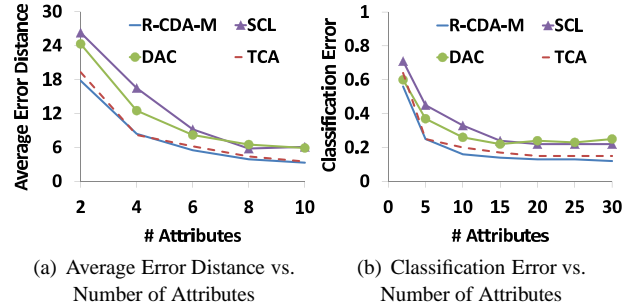


Figure 1: [Lower is better] (left) Average Error Distance on the Bike data and (right) Classification Error on the Energy data against the number of attributes in the mapping space. For regression, we use Regularized Least Square Regression. For classification, we use SVM.

sumed, and amount of heating. The second one (\mathbf{D}_Y) has 3257 records and 20 attributes; some of which are *electricity consumption*, *indoor CO2 concentration*, and *indoor air temperature*. Two data sets have 7 attributes in common. Further, they are collected in different time periods and hence there is no direct way to match their samples.

Regarding the results, RCDA_M identifies that attributes *number of staff members* and *percentage of building occupation* of \mathbf{X} are related to attribute *amount of heating* of \mathbf{Y} . Taking into account the fact that two buildings are close to each other and have similar architectural characteristics, this relation is intuitively understandable. RCDA_M also detects that attribute *amount of water consumed* of \mathbf{X} is related to attributes *indoor CO2 concentration* and *indoor air temperature* of \mathbf{Y} . Interestingly, a similar relation has recently been found in [26] for data from a *single* building. That RCDA_M discovers this relation for two different buildings could be attributed to their close proximity and architectural similarities.

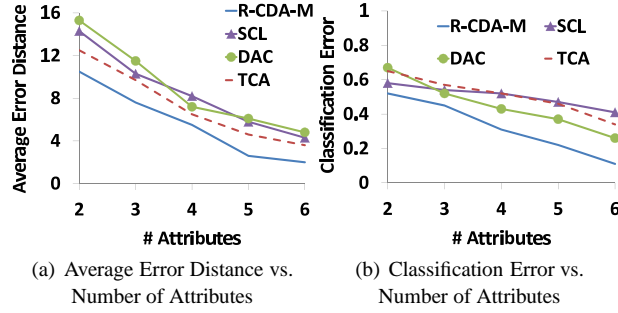


Figure 2: [Lower is better] (left) Average Error Distance and (right) Classification Error on the Building data against the number of attributes in the mapping space. For regression, we use Regularized Least Square Regression. For classification, we use SVM.

Besides, we repeat the cross-domain learning setup in Section 8.2. For regression, the target variable is *electricity consumption*. For classification, for each sample of both \mathbf{X} and \mathbf{Y} we choose the weekday when it was recorded as its class label. To make SCL, DAC, and TCA applicable, when testing with them we use the common attributes only. The results are in Figures 2(a) and 2(b). RCDA_M gives the best performance in both tasks. This could be attributed to the fact that RCDA_M can use all attributes to learn the mappings. The other three methods require $m = l$ and hence do not take full advantage of the available information for joint analysis of \mathbf{X} and \mathbf{Y} .

8.4 Non-redundant Subspace Clustering Next, we apply CDA to compare subspace clusters. Subspace clustering aims at finding objects clustered in subsets, *subspaces*, of all dimensions. There are two available approaches. One that finds subspace clusters that do not overlap in terms of objects [32] and one that does allow overlap in the objects [19]. We consider the latter setting. One of its major problems is redundancy in the output, as many clusters may be syntactically different but convey similar information. RESCU [24] is a recently proposed technique to alleviate the issue and improve post-analysis.

It works in an iterative manner. At each step, let \mathcal{C}_{sel} be the set of subspace clusters already found and \mathcal{C}_{cand} be the set of candidate subspace clusters. The potential of each cluster $C \in \mathcal{C}_{cand}$ w.r.t. \mathcal{C}_{sel} is defined as $\text{potent}(C, \mathcal{C}_{sel}) = \frac{|\text{cov}(C) \setminus \bigcup_{C' \in \mathcal{C}_{sel}} \text{cov}(C')|}{\text{cost}(C)}$ where $\text{cov}(C)$ is the set of objects covered by C and $\text{cost}(C)$ is the cost of mining C . RESCU then selects $C^* \in \mathcal{C}_{cand}$ with largest potential. The process goes on until \mathcal{C}_{cand} becomes empty. That RESCU considers object counts only is simplistic.




Figure 3: [Higher is better] Clustering quality (F1 values) of original RESCU and RESCU with RCDA_M on real-world data sets.

To improve, we propose to set

$$\text{potent}(C, \mathcal{C}_{sel}) = \frac{|\text{cov}(C) \setminus \bigcup_{C' \in \mathcal{C}_{sel}} \text{cov}(C')|}{\text{cost}(C) \sum_{C' \in \mathcal{C}_{sel}} \text{dist}(C, C')}$$

where $\text{dist}(C, C')$ is the distance between \mathbf{D}_C and $\mathbf{D}_{C'}$, respectively the data of C and C' . Note that \mathbf{D}_C and $\mathbf{D}_{C'}$ may have nonidentical sets of records and attributes. We define dist as follow. For each CDA method \mathcal{A} , we set

$$\text{dist}_{\mathcal{A}}(C, C') = \frac{1}{w} \sum_{i=1}^w \text{obj}_{\mathcal{A}}(u_i, v_i)$$

where w is the minimum of the dimensionality of \mathbf{D}_C and $\mathbf{D}_{C'}$, $\text{obj}_{\mathcal{A}}$ is the objective function of \mathcal{A} , and (u_i, v_i) is the i^{th} pair of canonical vectors identified by \mathcal{A} .

Therefore, we take into account also the similarity of C to the clusters in \mathcal{C}_{sel} . We hypothesize that this would improve the clustering quality. Note that existing methods for comparing clusters assume that clusters are on the same space [37] and hence are not applicable here.

For evaluation, we use the same UCI data sets employed in [24]. We report average F1 scores, the harmonic mean of recall and precision, over 10 independent runs. The results are in Figure 3. We find that RCDA_M consistently improves the clustering quality of the original RESCU for all data sets tested. This confirms both the applicability of CDA in reducing redundancy of subspace clustering, as well as our hypothesis. Further study will have to show to what extent CDA can be used further to discover interestingly different sub-parts of a data set.

9 Conclusion

We studied the problem of analyzing the relations of two random vectors of potentially different number of attributes, realizations, and possibly not even having a joint distribution. We proposed CDA, a formalization of this problem, and three practical instantiations. We also introduced two alternative formulations that provide a wider gamut to tackle the problem. Experimental results on various tasks evidence

the potential of our methods. For future work, we plan to look at non-linear transformations using kernel methods and deep learning. We are investigating how to measure privacy preservation and data security between original and released data using CDA.

Acknowledgements

The authors are supported by the Cluster of Excellence “Multimodal Computing and Interaction” within the Excellence Initiative of the German Federal Government.

References

- [1] T. Aho, B. Zenko, S. Dzeroski, and T. Elomaa. Multi-target regression with rule ensembles. *JMLR*, 13:2367–2407, 2012.
- [2] S.-I. Amari. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.
- [3] G. Andrew, R. Arora, J. Bilmes, and K. Livescu. Deep canonical correlation analysis. In *ICML (3)*, pages 1247–1255, 2013.
- [4] M. T. Bahadori, Y. Liu, and D. Zhang. Learning with minimum supervision: A general framework for transductive transfer learning. In *ICDM*, pages 61–70, 2011.
- [5] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy layer-wise training of deep networks. In *NIPS*, pages 153–160, 2006.
- [6] J. Blitzer, S. Kakade, and D. P. Foster. Domain adaptation with coupled subspaces. In *AISTATS*, pages 173–181, 2011.
- [7] J. Blitzer, R. T. McDonald, and F. Pereira. Domain adaptation with structural correspondence learning. In *EMNLP*, pages 120–128, 2006.
- [8] K. M. Borgwardt, A. Gretton, M. J. Rasch, H.-P. Kriegel, B. Schölkopf, and A. J. Smola. Integrating structured biological data by kernel maximum mean discrepancy. In *ISMB*, pages 49–57, 2006.
- [9] B. Chang, U. Krüger, R. Kustra, and J. Zhang. Canonical correlation analysis based on hilbert-schmidt independence criterion and centered kernel target alignment. In *ICML (2)*, pages 316–324, 2013.
- [10] R. Cilibrasi and P. M. Vitányi. Clustering by compression. *IEEE Trans. Inf. Theory*, 51(4):1523–1545, 2005.
- [11] W. Dai, Q. Yang, G.-R. Xue, and Y. Yu. Self-taught clustering. In *ICML*, pages 200–207, 2008.
- [12] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, 20(2):303–353, 1998.
- [13] A. Frank and A. Asuncion. UCI machine learning repository, 2010.
- [14] G. H. Golub and H. Zha. *The canonical correlations of matrix pairs and their numerical computation*. Springer, 1995.
- [15] A. Gretton, O. Bousquet, A. J. Smola, and B. Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *ALT*, pages 63–77, 2005.
- [16] D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural Computation*, 16(12):2639–2664, 2004.
- [17] H. Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936.
- [18] M. Karasuyama and M. Sugiyama. Canonical dependency analysis based on squared-loss mutual information. *Neural Networks*, 34:46–55, 2012.
- [19] H.-P. Kriegel, P. Kröger, and A. Zimek. Clustering high-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering. *TKDD*, 3(1), 2009.
- [20] A. Krishnamurthy, K. Kandasamy, B. Póczos, and L. A. Wasserman. On estimating l22 divergence. In *AISTATS*, 2015.
- [21] Q. V. Le, A. Karpenko, J. Ngiam, and A. Y. Ng. ICA with reconstruction cost for efficient overcomplete feature learning. In *NIPS*, pages 1017–1025, 2011.
- [22] E. Levina and P. J. Bickel. The earth mover’s distance is the mallows distance: Some insights from statistics. In *ICCV*, pages 251–256, 2001.
- [23] S. Liu, M. Yamada, N. Collier, and M. Sugiyama. Change-point detection in time-series data by relative density-ratio estimation. *Neural Networks*, 43:72–83, 2013.
- [24] E. Müller, I. Assent, S. Günnemann, R. Krieger, and T. Seidl. Relevant subspace clustering: Mining the most interesting non-redundant concepts in high dimensional data. In *ICDM*, pages 377–386, 2009.
- [25] H. V. Nguyen, E. Müller, and K. Böhm. 4S: Scalable subspace search scheme overcoming traditional apriori processing. In *IEEE BigData*, 2013.
- [26] H. V. Nguyen, E. Müller, J. Vreeken, P. Efron, and K. Böhm. Multivariate maximal correlation analysis. In *ICML*, pages 775–783, 2014.
- [27] J. Nocedal and S. Wright. *Numerical optimization*. Springer, 1999.
- [28] B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609, 1996.
- [29] M. Oresic, V. A. Hänninen, and A. Vidal-Puig. Lipidomics: a new window to biomedical frontiers. *Trends in Biotechnology*, 26(12):647–652, 2011.
- [30] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang. Domain adaptation via transfer component analysis. In *IJCAI*, pages 1187–1192, 2009.
- [31] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.*, 22(10):1345–1359, 2010.
- [32] D. Park, C. Caramanis, and S. Sanghavi. Greedy subspace clustering. In *NIPS*, pages 2753–2761, 2014.
- [33] K. Pearson. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *Philosophical Magazine*, 50:157–175, 1900.
- [34] S. T. Rachev. The monge-kantorovich mass transference problem and its stochastic applications. *Theory Probab. Appl.*, 29(4):647–676, 1984.
- [35] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng. Self-taught learning: transfer learning from unlabeled data. In *ICML*, pages 759–766, 2007.
- [36] M. Rao, S. Seth, J.-W. Xu, Y. Chen, H. Tagare, and J. C.

Príncipe. A test of independence based on a generalized correlation function. *Signal Processing*, 91(1):15–27, 2011.

- [37] S. Romano, J. Bailey, X. V. Nguyen, and K. Verspoor. Standardized mutual information for clustering comparisons: One step further in adjustment for chance. In *ICML*, pages 1143–1151, 2014.
- [38] M. Schweiker and M. Shukuya. Study on the effect of preference of air-conditioning usage on the exergy consumption pattern within a built environment. *International Journal of Exergy*, 11(4):409–422, 2012.
- [39] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, London, 1986.
- [40] K. Sridharan and S. M. Kakade. An information theoretic framework for multi-view learning. In *COLT*, pages 403–414, 2008.
- [41] N. Tatti and J. Vreeken. Comparing apples and oranges: measuring differences between exploratory data mining results. *Data Min. Knowl. Discov.*, 25(2):173–207, 2012.
- [42] A. L. Traud, E. D. Kelsic, P. J. Mucha, and M. A. Porter. Comparing community structure to characteristics in online collegiate social networks. *SIAM Review*, 53(3):526–543, 2011.
- [43] A. Tripathi, A. Klami, M. Oresic, and S. Kaski. Matching samples of multiple views. *Data Min. Knowl. Discov.*, 23(2):300–321, 2011.
- [44] A. Wagner, T. Lützkendorf, K. Voss, G. Spars, A. Maas, and S. Herkel. Performance analysis of commercial buildings—results and experiences from the german demonstration program ‘energy optimized building (EnOB)’. *Energy and Buildings*, 68:634–638, 2014.
- [45] X. Yin. Canonical correlation analysis based on information theory. *Journal of Multivariate Analysis*, 91(2):161–176, 2004.
- [46] M. Zhang, M. Hadjieleftheriou, B. C. Ooi, C. M. Procopiuc, and D. Srivastava. Automatic discovery of attributes in relational databases. In *SIGMOD Conference*, pages 109–120, 2011.

A Proofs

Proof. (Lemma 4.1) We prove for the case when $n \neq k$. The proof for when $n = k$ follows similarly. W.l.o.g., we assume that $n < k$.

As aforementioned, since $n \neq k$, we replicate each realization such that both distributions have $n \times k$ realizations. We write the new sets of realizations as $\{x'_1, \dots, x'_{n \times k}\}$ and $\{y'_1, \dots, y'_{n \times k}\}$. We have

$$(A.1) \quad \text{diff}(p(X), q(Y)) \sim \left(\sum_{\{j_1, \dots, j_{n \times k}\}} \sum_{i=1}^{n \times k} |x'_i - y'_{j_i}|^t \right)^{1/t}$$

where $\{j_1, \dots, j_{n \times k}\}$ is a permutation of $\{1, \dots, n \times k\}$. For each $i \in [1, n]$, and each $j \in [1, k]$, the term $|x_i - y_j|^t$ appears at most $\min(n, k) = n$ in a single permutation. Thus, the total number of times $|x_i - y_j|^t$ appears in the sum is

$$\sum_{c=1}^n c(n \times k - n)! \prod_{r=0}^{n-c-1} (n \times k - k - r).$$

As we make no preference in picking i and j , every term $|x_i - y_j|^k$ appears the same number of times as specified above. Hence, we arrive at the result.

Proof. (Lemma 5.1) From $u^T u = v^T v = 1$, it holds that $|u^T \mathbf{X}| \leq \sqrt{m}$ and $|v^T \mathbf{Y}| \leq \sqrt{l}$. Therefore, with $\beta = \sqrt{\frac{m}{l}}$, we have that $\beta v^T \mathbf{Y} \leq \sqrt{m}$.

Proof. (Lemma 5.2) Let $u' \in \mathbb{R}^{m+c}$ be such that $u' = [u, \mathbf{0}]$ and $v' \in \mathbb{R}^{l+d}$ be such that $v' = [v, \mathbf{0}]$. We have that $(u')^T u' = (v')^T v' = 1$ and

$$\text{diff}(p((u')^T \mathbf{X}), q(\beta(v')^T \mathbf{Y})) = \text{diff}(p(u^T \mathbf{X}), q(\beta v^T \mathbf{Y})).$$

That is, $\text{diff}(p((u')^T \mathbf{X}), q(\beta(v')^T \mathbf{Y}))$ is small, i.e. (u', v') and hence (u, v) are identifiable by CDA.

Proof. (Proposition 6.1) Our proof is based on [21]. In particular, we have

$$\begin{aligned} & \frac{\lambda}{n} \sum_{i=1}^n \|u u^T \tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_i\|^2 \\ &= \lambda \times \text{tr} \left[(u u^T - \mathbf{I})^T (u u^T - \mathbf{I}) \times \frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{x}}_i (\tilde{\mathbf{x}}_i)^T \right] \end{aligned}$$

As \mathbf{X} is whitened, $\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{x}}_i (\tilde{\mathbf{x}}_i)^T = \mathbf{I}$. Hence,

$$\frac{\lambda}{n} \sum_{i=1}^n \|u u^T \tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_i\|^2 = \|u u^T - \mathbf{I}\|_F^2.$$

Proof. (**Proposition 6.2**) Our proof again is based on [21]. In particular, we have

$$\begin{aligned}
& (u^T u - 1)^2 \\
&= (u^T u - 1)(u^T u - 1) \\
&= u^T u u^T u - 2u^T u + 1 \\
&= \text{tr}(u u^T u u^T) - 2\text{tr}(u u^T) + \text{tr}(\mathbf{I}_m) + 1 - m \\
&= \text{tr}[u u^T u u^T - 2u u^T + \mathbf{I}_m] + 1 - m \\
&= \text{tr}[(u u^T - \mathbf{I}_m)^T (u u^T - \mathbf{I}_m)] + 1 - m \\
&= \|u u^T - \mathbf{I}_m\|_F^2 + 1 - m.
\end{aligned}$$

Proof. (**Lemma 6.1**) From Propositions 6.1 and 6.2, we can see that when \mathbf{X} and \mathbf{Y} are whitened, RCDA is equivalent to

$$\arg \min_{\substack{u^T u^j = 0, \forall j \in [1, i-1] \\ v^T v^j = 0, \forall j \in [1, i-1]}} [\lambda(u^T u - 1)^2 + \delta(v^T v - 1)^2 + \text{diff}(p(u^T \mathbf{X}), q(\beta v^T \mathbf{Y}))].$$

When $\lambda \uparrow +\infty$ and $\delta \uparrow +\infty$, the unit norm costs reduce to hard constraints.

B Sample Estimator of the Quadratic Measure

We have $\text{diff}(p(Z), q(Z))$ is equal to

$$\begin{aligned}
\text{diff}(p(Z), q(Z)) &= \int p(z)p(z)dz + \int q(z)q(z)dz \\
&\quad - \int q(z)p(z)dz - \int p(z)q(z)dz
\end{aligned}$$

which is $E_p(p(Z)) + E_q(q(Z)) - E_p(q(Z)) - E_q(p(Z))$.

With KDE, $\text{diff}(p(Z), q(Z))$ becomes

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n \hat{p}_h(x_i) + \frac{1}{k} \sum_{j=1}^k \hat{q}_b(y_j) \\
& - \frac{1}{n} \sum_{i=1}^n \hat{q}_b(x_i) - \frac{1}{k} \sum_{j=1}^k \hat{p}_h(y_j)
\end{aligned}$$

Thus, we arrive at the result.

C Multi-dimensional Reconstruction CDA

We form this alternative formulation by combining MCDA and RCDA. Formally, it is defined as follow.

DEFINITION 4. *Multi-dimensional RCDA (MRCDA) solves for $\mathbf{U} = (u^1, \dots, u^r) \in \mathbb{R}^{m \times r}$ and $\mathbf{V} = (v^1, \dots, v^r) \in \mathbb{R}^{l \times r}$ where (\mathbf{U}, \mathbf{V}) is the solution of*

$$\arg \min_{\mathbf{U}^T \mathbf{U} = \mathbf{V}^T \mathbf{V} = \mathbf{I}_r} \left(\frac{\lambda}{n} \sum_{i=1}^n \|\mathbf{U} \mathbf{U}^T \tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_i\|^2 + \frac{\delta}{k} \sum_{j=1}^k \|\mathbf{V} \mathbf{V}^T \tilde{\mathbf{y}}_j - \tilde{\mathbf{y}}_j\|^2 + \text{diff}(p(\mathbf{U}^T \mathbf{X}), q(\mathbf{V}^T \mathbf{Y})) \right)$$

with $\tilde{\mathbf{x}}_i = \mathbf{x}_i - \mu(\mathbf{X}_i)$ being the centered version of \mathbf{x}_i , $\tilde{\mathbf{y}}_j = \mathbf{y}_j - \mu(\mathbf{Y}_j)$ being the centered version of \mathbf{y}_j , $p(\cdot)$ and $q(\cdot)$ being pdfs, diff being a divergence measure of pdfs, $\lambda > 0$ and $\delta > 0$ being weights of reconstruction costs, and $\mathbf{\Gamma} = \text{diag}(\beta_1, \dots, \beta_r)$ being the scaling matrix to bring $(u^i)^T \mathbf{X}$ and $(v^i)^T \mathbf{Y}$ to the same domain.

To solve MRCDA, we again use PE divergence. Note that for both MCDA and MRCDA, the bandwidths are: $\sigma_{\mathbf{X}} = r \times \text{median}\{\|\mathbf{x}_i - \mathbf{x}_j\| : 1 \leq i, j \leq n\}$ and $\sigma_{\mathbf{Y}} = (\sum_{k=1}^r \beta_i) \times \text{median}\{\|\mathbf{y}_i - \mathbf{y}_j\| : 1 \leq i, j \leq k\}$.

D Sensitivity to Noisy Attributes

To assess quality of all methods tested to noisy attributes, we extend our setup in Section 8.1. In particular, we first generate n samples for $\mathbf{X} \in \mathbb{R}^{5+c}$ where $c \geq 2$ and $\mathbf{X}_i \sim \mathcal{N}(0, 1)$. Then we generate k samples for $\mathbf{Y} \in \mathbb{R}^{3+c}$ where \mathbf{Y}_1 to \mathbf{Y}_4 follow the three types of relations: linear, mixed, and non-linear; in addition, $\mathbf{Y}_j \sim \mathcal{N}(0, 1)$ for $j \in [5, 3+c]$. Hence, c controls the number of noisy attributes. For exposition, in Table 2 we present results on the setting: $n = k$, row order intact, and non-linear relations. We can see that our methods are relatively stable to noisy dimensions, which is in line with our theoretical analysis in Section 5. This highlights the benefits of our setting of β .

E Optimizing vs. Not Optimizing β

We use the setup in Section 8.1 to compare CDA methods with and without optimizing β . The results are in Table 3. Overall, we can see that there is no clear winner between the two options. This could be explained as follows. When β is a variable to be optimized, the search space and hence the optimization problem become more complex, especially for CDA_Q and CDA_P . With a complex search space the optimizers employed may be more prone to local optimum. This explains why optimizing β currently does not bring a clear improvement, while costing more time. In particular, optimizing β incurs from 8% to 15% more time than iteratively setting it; please see the last line of Table 3 for more details.

F Efficiency

Above we show that iteratively setting β gives better efficiency than optimizing it. There is another factor that affects the runtime, which is the reconstruction costs. To study this, we use the runtime obtained from our experiments in Section 8.1. We summarize the results in Table 4. Overall, we observe that using reconstruction costs bring about 5 times speedup compared to the other variants. This could be attributed to the fact that CDA with reconstruction costs yields unconstrained optimization problems where fast unconstrained optimizer, like L-BFGS, can be applied.

c	CDA			CCA	LCCA	SCL	DAC	SSC
	CDA _Q	RCDA _M	MRCDA _P					
2	0.36	0.35	0.35	0.50	0.48	0.52	0.57	0.54
4	0.38	0.36	0.36	0.53	0.51	0.55	0.59	0.56
6	0.35	0.38	0.36	0.53	0.53	0.53	0.63	0.59
8	0.38	0.36	0.37	0.56	0.56	0.56	0.66	0.60
10	0.39	0.38	0.39	0.58	0.56	0.58	0.67	0.61

Table 2: [Lower is better] Frobenius norm errors against the number of noisy attributes for the setting: $n = k$, row order intact, and non-linear relations. All values are obtained after 10 trials. Best values and comparable ones according to the Wilcoxon signed rank test at $\alpha = 5\%$ are in **bold**.

Setting	Type	Iteratively Setting β			Optimizing β		
		CDA _Q	RCDA _M	MRCDA _P	*CDA _Q	*RCDA _M	*MRCDA _P
$n = k$	Linear	0.31	0.30	0.31	0.33	0.29	0.32
	Mixed	0.34	0.34	0.35	0.35	0.36	0.34
	Non-linear	0.36	0.35	0.35	0.34	0.35	0.34
$n \neq k$	Linear	0.33	0.31	0.30	0.33	0.34	0.31
	Mixed	0.36	0.36	0.38	0.35	0.37	0.37
	Non-linear	0.39	0.38	0.39	0.40	0.41	0.40
Average runtime		1473	217	284	1611	252	319

Table 3: [Lower is better] Frobenius norm errors for iteratively setting β vs. optimizing it. All values are obtained after 10 trials. Best values and comparable ones according to the Wilcoxon signed rank test at $\alpha = 5\%$ are in **bold**. The superscript “*” indicates CDA methods that optimize β . The last line is the average runtime (in seconds) of all methods.

No reconstruction costs			Reconstruction costs		
CDA _Q	CDA _M	MCDA _P	RCDA _Q	RCDA _M	MRCDA _P
1473	1265	1622	245	217	284

Table 4: [Lower is better] Average runtime (in seconds) of CDA methods not using reconstruction costs and the respective ones using reconstruction costs.